



US006978389B2

(12) **United States Patent**
Jahnke

(10) **Patent No.:** **US 6,978,389 B2**
(45) **Date of Patent:** **Dec. 20, 2005**

(54) **VARIABLE CLOCKING IN AN EMBEDDED SYMMETRIC MULTIPROCESSOR SYSTEM**

(75) Inventor: **Steven R. Jahnke**, Tokyo (JP)
(73) Assignee: **Texas Instruments Incorporated**,
Dallas, TX (US)
(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 516 days.

(21) Appl. No.: **10/256,698**

(22) Filed: **Sep. 27, 2002**

(65) **Prior Publication Data**
US 2003/0120963 A1 Jun. 26, 2003

Related U.S. Application Data
(60) Provisional application No. 60/342,911, filed on Dec. 20, 2001.
(51) **Int. Cl.**⁷ **G06F 1/16**
(52) **U.S. Cl.** **713/322; 713/300; 713/375; 713/401**

(58) **Field of Search** **713/300, 322, 713/375, 401**

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,812,860 A *	9/1998	Horden et al.	713/322
6,141,762 A *	10/2000	Nicol et al.	713/300
6,578,123 B1 *	6/2003	Austin et al.	711/165
6,664,775 B1 *	12/2003	Clark et al.	323/351
6,675,304 B1 *	1/2004	Pole et al.	713/322

* cited by examiner

Primary Examiner—Thomas Lee
Assistant Examiner—Vincent Tran
(74) *Attorney, Agent, or Firm*—Robert D. Marshall, Jr.; W. James Brady, III; Frederick J. Telecky, Jr.

(57) **ABSTRACT**

In multiprocessor systems the task of holding power dissipation to its lowest possible level is challenging. This invention permits reduced power dissipation by optionally independently clocking selected central processing units at lower frequencies if they are not fully loaded. The variable clocking system enables synchronization between central processing units operating a differing frequencies and shared memory and peripherals. This allows for significant power reduction in the frequently occurring scenario where all processors are not driven to their limits by prevailing system requirements.

6 Claims, 4 Drawing Sheets

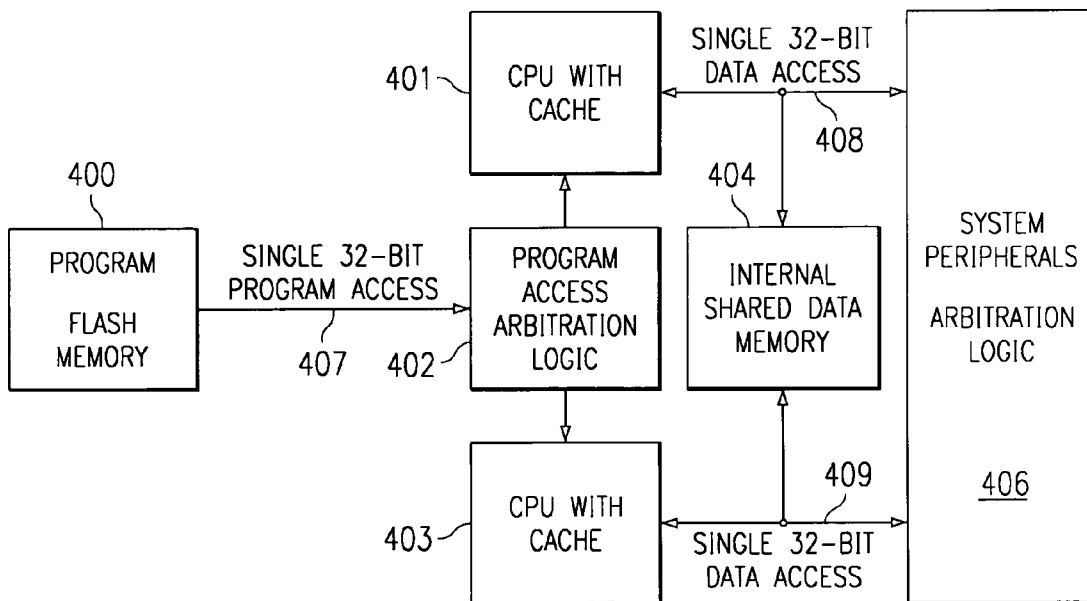


FIG. 1
(PRIOR ART)

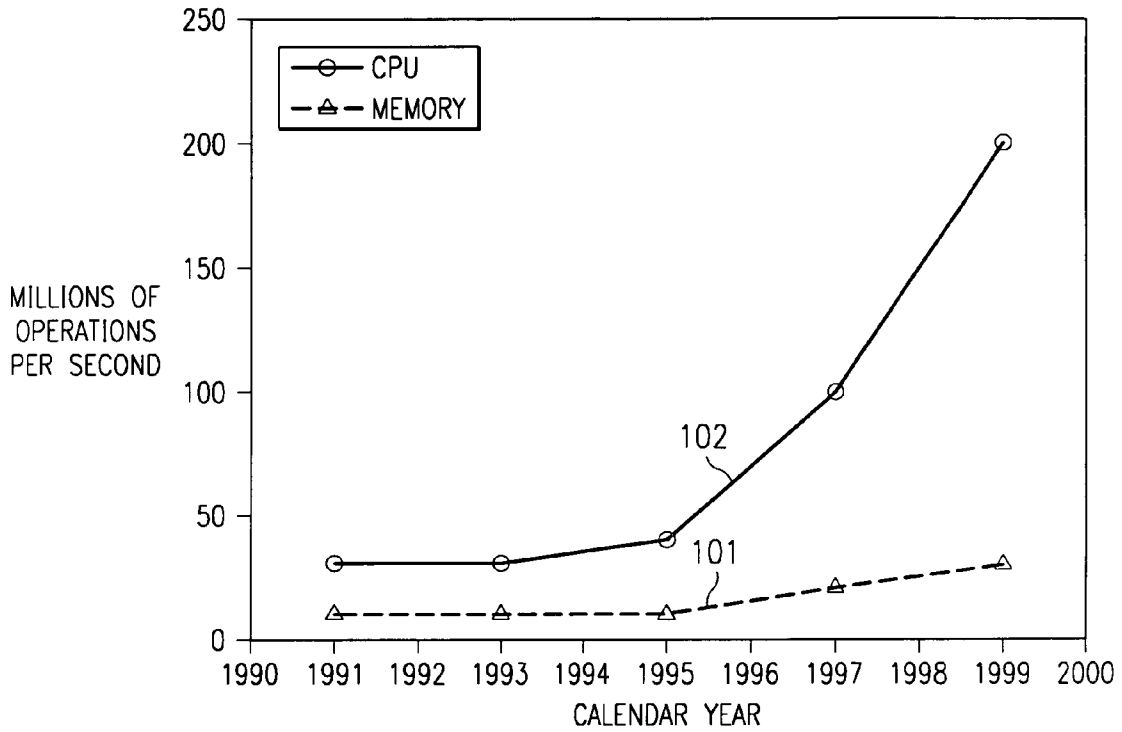
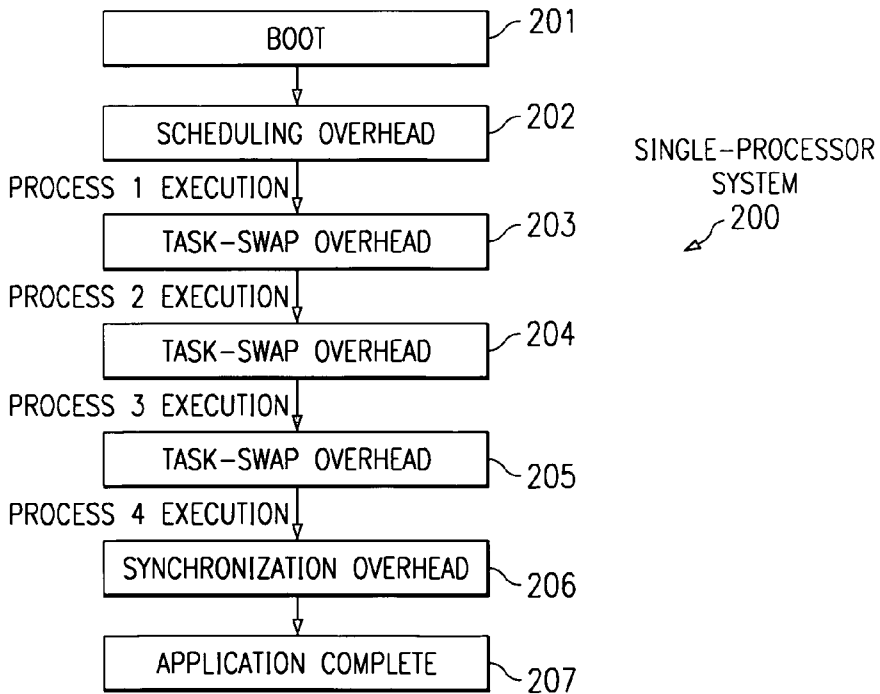


FIG. 2
(PRIOR ART)



MULTI-PROCESSOR SYSTEM
300

FIG. 3
(PRIOR ART)

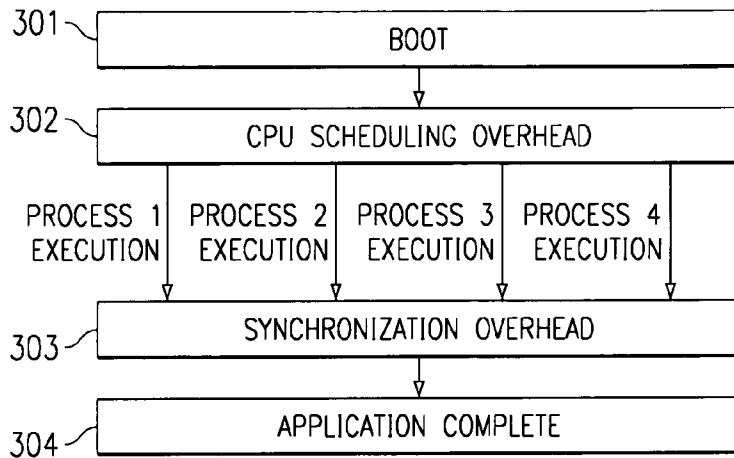


FIG. 4

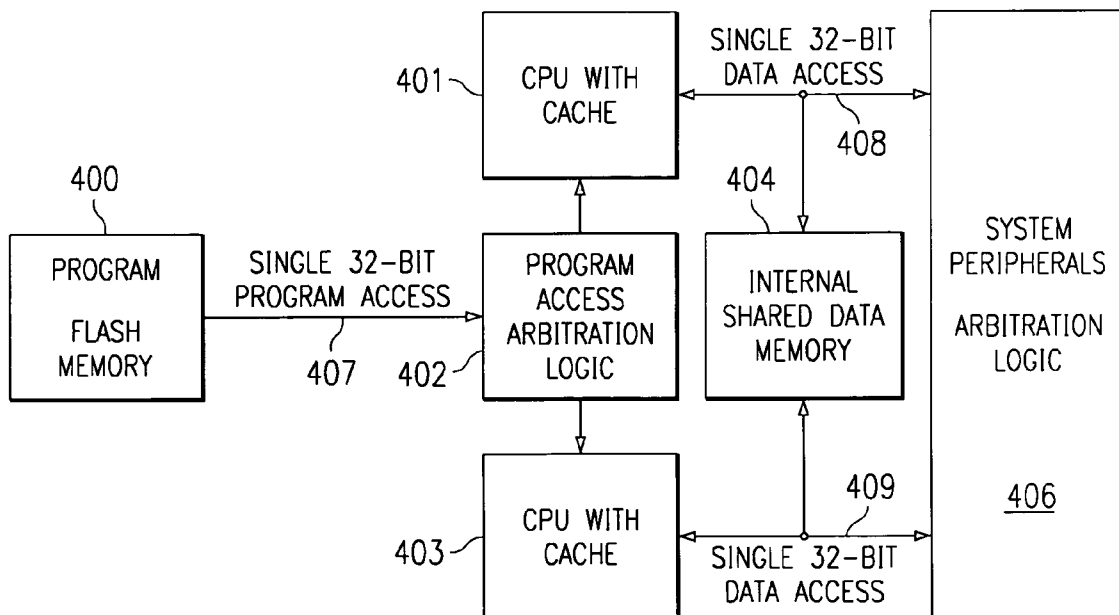


FIG. 5
(PRIOR ART)

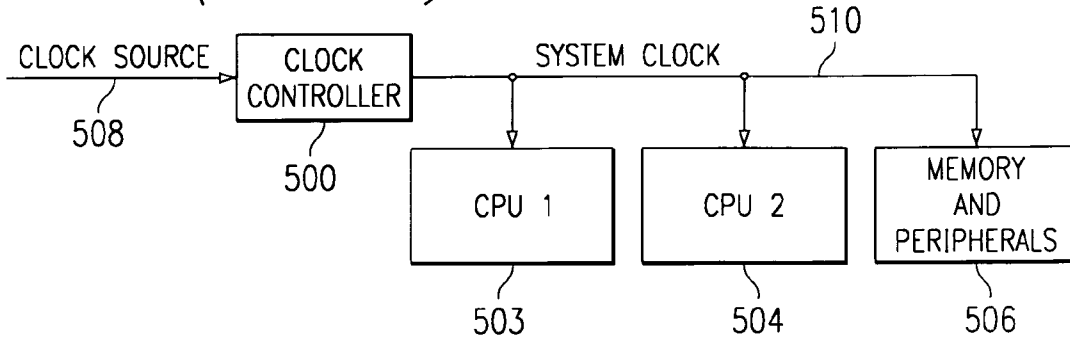
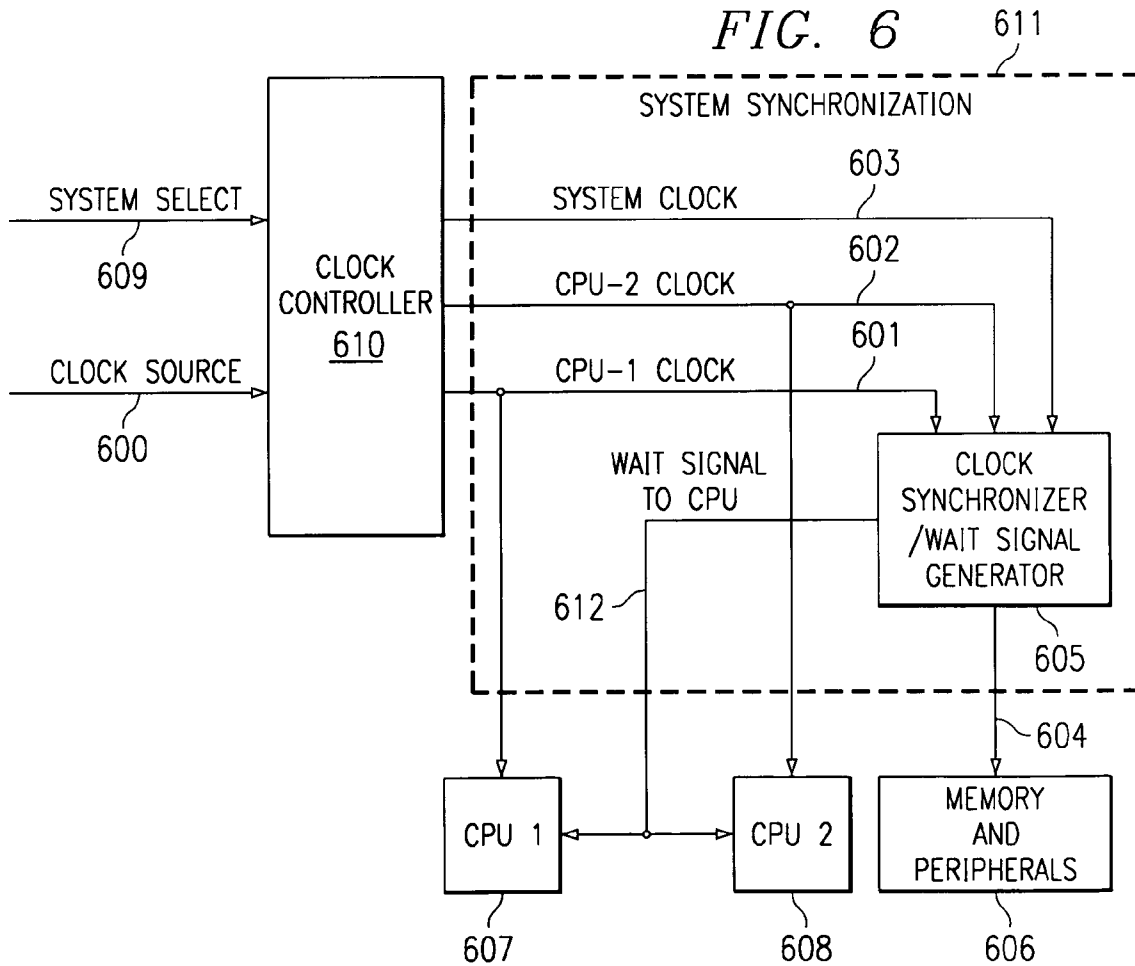
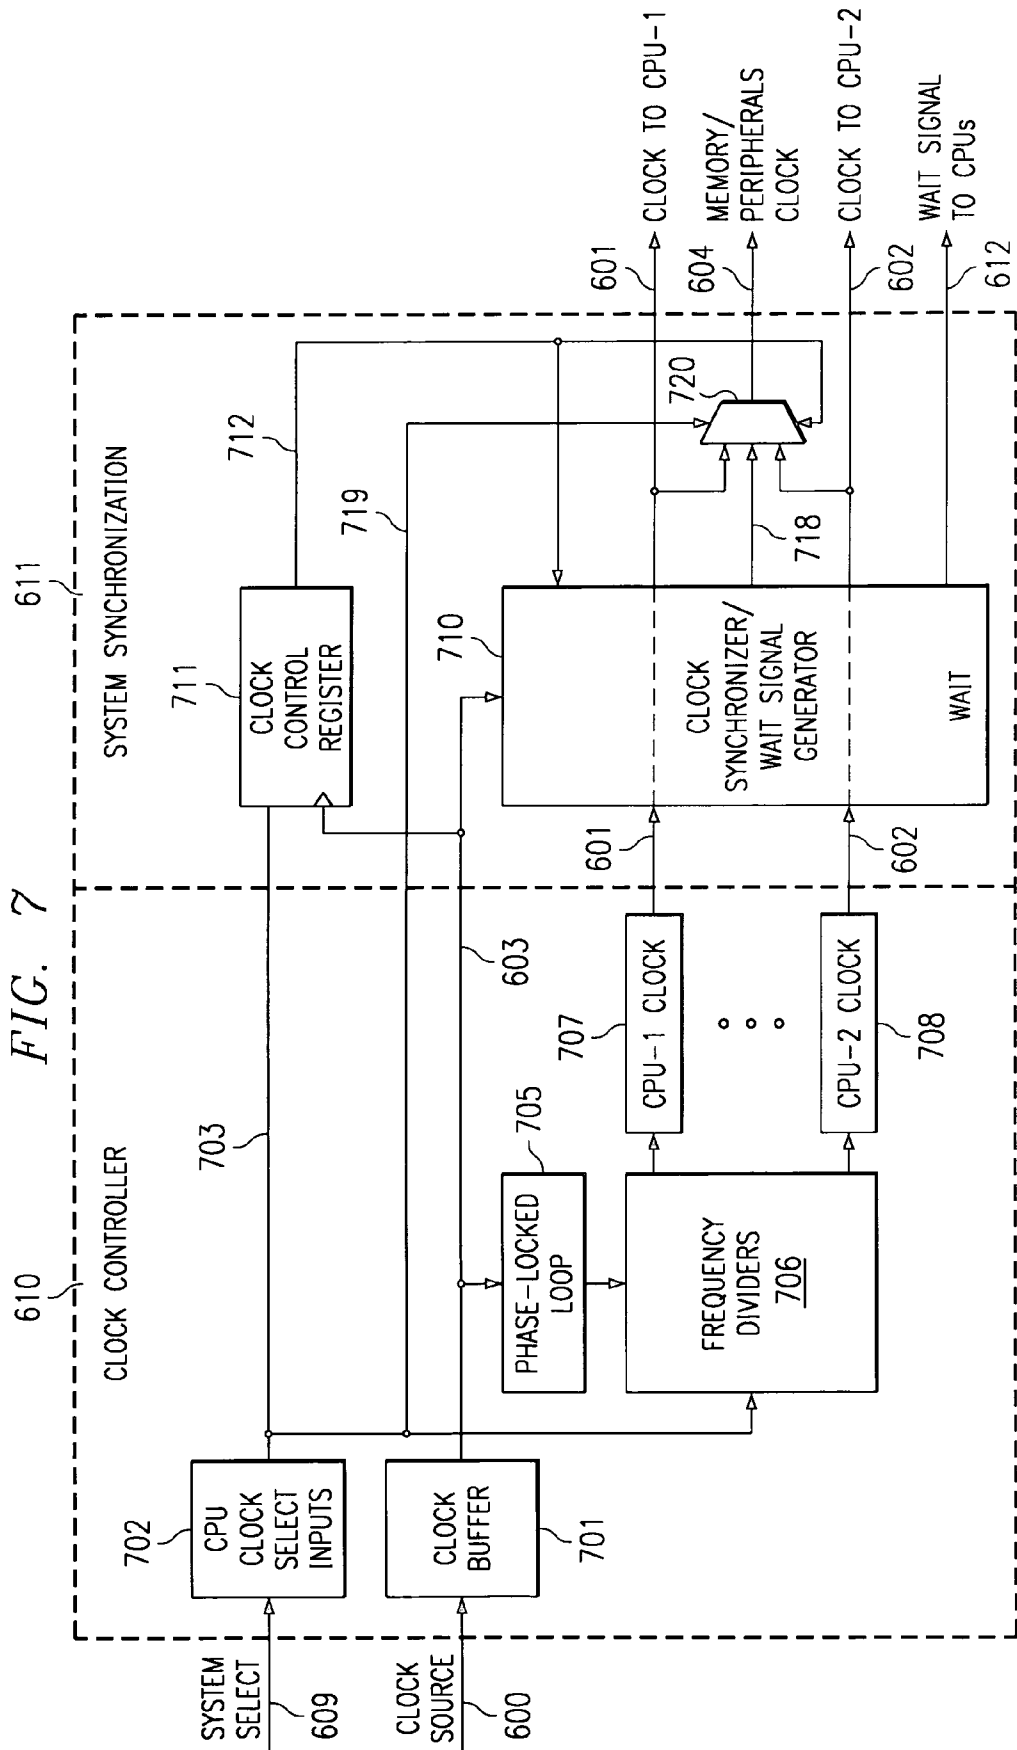


FIG. 6





VARIABLE CLOCKING IN AN EMBEDDED SYMMETRIC MULTIPROCESSOR SYSTEM

This application claims priority under 35 USC §119(e)(1) of Provisional Application No. 60/342,911, filed Dec. 20, 2001.

TECHNICAL FIELD OF THE INVENTION

The technical field of this invention is computer systems and more particularly multiprocessor computer systems.

BACKGROUND OF THE INVENTION

As each generation of silicon process technology has provided increasing integration density using smaller geometry transistors, central processing unit architects have continually debated how to use the additional device area to increase application performance. With smaller lower capacitance transistors, operating frequency has proportionally increased, yielding a direct performance gain. However, the access time of the memory function that holds the application program has not kept pace with the speed increases in the central processing unit. This is illustrated in FIG. 1. Memory speed improvement **101** has been gradual. Central processing unit speed improvement **102** has been more marked.

As a result, the performance gain that should be realizable from central processing unit operating frequency advances cannot be achieved without corresponding architectural enhancements in the central processing unit program memory path. As noted in FIG. 1, the speed difference between memory and processors has greatly increased in the past few years. As this gap continues to grow, the memory central processing unit interface will have an even greater effect on overall system performance. The traditional solution to reduce the effect of the central processing unit memory interface bottleneck is to use some form of memory hierarchy. In a general-purpose application processor, a cache system is employed that will allow the hardware at run time to keep copies of the most commonly used program elements in faster, internal RAM. In a more deeply embedded, performance sensitive application (such as a DSP), a form of tightly coupled memory is used that will allow the software to copy either a part of or all of the application program into on-chip RAM. In both of these techniques, the hardware architect gains system performance by the direct, brute force method of simply increasing clock frequency. This solution has proven successful because the performance gains by process technology alone have proved enough for current embedded applications, and there is no impact on application developers to migrate to a faster higher performance system.

It is important, for the clear exposition of processor techniques that follow, to define first the term embedded processor system (EPS) as employed here and as differentiated from a conventional non-embedded multi-chip processor system (MCPS). An embedded processor system includes a processor system integrated on a single chip having one or more central processing units plus a full complement of functional features and functional elements. This full complement of features, not normally included in conventional non-embedded multi-chip processor systems (MCPS). The MCPS is formed from one or more single chip central processing units and additional packaged devices performing memory, interface and peripheral circuits and these are assembled on a printed-wire board (PWB).

Additionally we define the embedded multiprocessor system (EMPS) as having multiple central processing units, complex memory architectures and a wide range of peripheral devices all fully integrated on a single chip. Such a system normally includes another special peripheral, an external memory interface (EMIF) coupled to a large amount of external memory. Central processing unit interactions and cache interactions on an embedded processor clearly involve more complex functionality when compared to a non-embedded processor device. Further, the embedded multiprocessor is typically used in a real-time environment leading to additional requirements for the coherent handling of interrupt operations and power consumption control.

The design methodologies used to support existing processors create a bottleneck in the ability for central processing unit designers to maximize frequency gain without extraordinary effort. At the same time the type of applications being considered for next generation embedded processors grows significantly in complexity. Application performance demand outpaces the ability of designers to efficiently provide performance through operating frequency alone at a reasonable development cost.

The disparity between embedded processor application performance requirements and performance gain through operating frequency alone has not gone unnoticed. In many new digital signal processors, two distinct paths have been used to affect increased system performance. The first technique is the use of enhanced central processing unit architectures having instruction level parallelism and the second technique is the use of system task specialization among different types of simpler but more specialized processors. These two paths are outlined below.

The Texas Instruments TMS320C6000 family of digital signal processors provides an example demonstrating the use of an effective central processing unit architecture to gain performance. Many of these devices use a form of instruction level parallelism (ILP) called very long instruction word (VLIW) to extract a performance gain by analyzing the code behavior at the most basic instruction level. The compiler effectively schedules unrelated instructions to be executed in two or more parallel processing units. This allows the processor to do work on more than one instruction per cycle. Since the instruction scheduling and analysis is done by the compiler, the hardware architecture can be simplified somewhat over other forms of instruction level parallelism ILP, such as super-scalar architectures.

Due to this emphasis on the compiler-based performance extraction, there is little impact on the task of application programmers. Application development can be done in a high-level language and be compiled normally. This is done in a non-ILP based system. This ease of application development, coupled with a performance gain without an operating frequency increase has resulted in the success of this form of enhancement. However, these benefits do not come without cost. Both the development effort in creating a new instruction set architecture (ISA), along with the compiler optimizations required are significant. In the future, once the underlying architecture is fixed, the only means of gaining additional performance is by increasing operating frequency.

Other Texas Instruments digital signal processors, the so-called OMAP devices and the TMS320C5441 provide examples of the technique of breaking the target application into fundamental domains and targeting a simpler processor to each domain. Based on system analysis, the system architect breaks the total application into smaller parts and puts together a separate programming plan for each central processing unit in place. In the past, this could have been

done only at the board level, where a specialized processor would be targeted for a specific application task. However, the integration density offered by current process enhancements allows these specialized central processing units to be placed on a single die. This enables a tighter coupling between the processors. Fundamentally, the application developer writes code as if he or she was dealing with each processor as an independent platform.

The programmer must be cognizant of the hardware architecture and program each processor independently. Greater coupling between the integrated processors allows for a more efficient passing of data than at the board level. However, the application is primarily written with the focus on the separate processors in the system. Code reuse and porting is difficult even among the processors in the same system, because each processor is really the centerpiece of its subsystem. Each processor may have a different memory map, different peripheral set and perhaps even a different instruction set (such as OMAP). In applications that have very distinct boundaries, such as a cell phone, this method of extracting performance is unparalleled. Each part of the application can be targeted to an optimized processor and programmed independently.

Development efforts are reduced somewhat since a new instruction set is not required to gain performance. However, from an application development and road map perspective, this technique does not offer the ease of use that instruction level parallelism offers. In many applications, there is no clear line where to divide the work. Even when done, the system cannot easily use all the performance of each central processing unit. If one central processing unit is idle while another is very busy, it is difficult to readjust central processing unit loading once the code has been written. If tighter coupling between the system processors is desired, significant software overhead must be added to insure data integrity.

SUMMARY OF THE INVENTION

In an embedded symmetric multiprocessor (ESMP) system, all devices are normally driven by the same clock to aid in maintaining data coherency and software synchronization. A basic premise of ESMP is that there is a symmetry among the processors in the system. The processors are of the same type and run at the same frequency. They are typically clocked by clocking networks that tightly control clock edge skew as seen at the clock input of the various processors.

Nevertheless, in both the conventional symmetric embedded multiprocessors (CSMP) of desktop systems and in embedded systems in general, there is a significant concern about power dissipation. Therefore, it is desirable to have selected system processors run at a lower frequency if they are not fully loaded. This invention describes a variable clocking system that allows mixed frequency operation.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other aspects of this invention are illustrated in the drawings, in which:

FIG. 1 illustrates the progress in speed performance of memory and central processor units in recent years according to the prior art;

FIG. 2 illustrates the execution time of plural processes by single-processor in accordance with the prior art;

FIG. 3 illustrates the execution time of plural processes by a multi-processor systems in accordance with the prior art;

FIG. 4 illustrates an example embedded symmetric multiprocessing system to which the invention is applicable;

FIG. 5 illustrates a conventional symmetric multiprocessor CSMP clocking system (Prior Art);

FIG. 6 illustrates a high-level block diagram of the variable clocking system of this invention; and

FIG. 7 illustrates in block diagram form the method of lower frequency clock generation and switching within the clock controller of the variable clocking system of this invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

The embedded symmetric multiprocessor system (ESMP) of this invention includes a powerful set of central processing unit-memory-peripheral functions densely integrated at the chip level. While some common multi-chip symmetric multiprocessor systems (MCSMP) are generally available at the board level now, the designer of MCSMP systems typically employs plural standard third or fourth generation central processing unit chips for the base processors. Two or more of these on standard processors are disposed on a mother-board and then connected by way of a commonly available bus interface device to a separate traffic controller and memory controller. Such systems use discrete interface and controller components and central processing unit-memory architectures at the board level. This combination of devices has a set of system interconnect requirements and concerns completely different from and often more troublesome than the embedded symmetric multiprocessor system (ESMP) of this invention. This invention has all interconnects of major importance implemented compactly at the chip level.

Conventional symmetric processor systems (CSMP) have been designed employing having two or more central processing units on the same die, cache architectures that include local unshared L1 caches for each central processing unit and generally an L2 cache shared by both central processing units. However, few if any conventional CSMP systems are available that have both the kind of primary memory normally present on the hard disc drive of a conventional desktop computer and include the full range of peripheral devices. We are labeling these conventional symmetric processor systems (CSMP) as non-embedded as they most often have simply a central processing unit-memory architecture with no peripheral system components. Processor designers have only recently initiated experimentation and research in the area of these higher complexity systems that include the full range of peripheral devices.

An embedded symmetric processor system (ESMP) includes the full complement of functional features and functional elements, such as peripheral functions and external memory interface, not contained in a conventional CSMP system. It contains integrated multiple central processing units with high complexity memory architectures plus peripherals (i.e. EDMA, UART, USB functions), plus local system memory and perhaps an interface to external memory if a large amount of memory is required. The central processing unit interactions and cache interactions on an embedded central processing unit are similar but much more complex than the CSMP case. These interactions must comprehend the cache interfaces with on-chip system memory and peripheral interfacing. Since the embedded processor is often used in a real-time environment, interrupt operations and the manner of controlling and reducing power consumption are handled differently.

5

In summary, the basic difference between the conventional symmetric multiprocessor (CSMP) and the embedded symmetric multiprocessor (ESMP), is that the conventional CSMP is simply a computing processor while the embedded symmetric multiprocessor ESMP is an integrated system having one or more central processing units plus a full complement of peripherals. A non-embedded CSMP deals with a simpler central processing unit-cache interface with minimal concerns for what happens beyond. An embedded ESMP must resolve a greater complexity of system interactions and interfaces requirements.

Both single-processor, instruction-level parallelism ILP architectures and the multi-processor architectures increase system performance by taking advantage of parallelism but at different levels of abstraction. A single processor ILP architectures takes advantage of unrelated central processing unit instructions that can be executed concurrently. The multiprocessor approach takes this a step further and makes use of unrelated application fragments that can be run concurrently. The instruction-level parallelism (ILP) approach has a very tight coupling between parallel operating units (i.e. execution units inside the core) and is completely hardware and software controlled. As a result, the parallelism is invisible to the user. The multiprocessor approach has very loose coupling between parallel operating units (i.e. separate processors) with minimal hardware and software interference. As a result, the parallelism is not only visible to the user, but system performance gain is dependent upon efficient division of the application. In applications that run more than one software process, there resides another level of parallelism in-between these two extremes: process level parallelism.

This invention includes a software process level that seeks system performance gain in process level parallelism using multiple central processing units. When a program first boots, the kernel, which may be either part of a real time operating system (RTOS) or custom-developed boot code, will schedule which parts of the application will be executed at which time. Some processes are launched based on a conditional event, such as the completion of a previous process or external event. However most major processes have some degree of independence from one another in a multi-tasking environment. The supervisor code from either the operating system or the boot-code schedules central processing unit time for each process, based on its priority. It is at this central processing unit scheduling point that additional performance can be gained through the use of multiple central processing units.

Instead of time-sharing all processes on a single central processing unit, the supervisor can split these processes among two or more central processing units. FIGS. 2 and 3 illustrate these two alternatives. FIG. 2 illustrates an example of the execution time for a complex process running on a single central processing unit system. FIG. 3 illustrates that same process running on a multiple central processing unit system. In these examples, four separate single processes are running.

On the single central processing unit system 200, each process is time shared on the single central processing unit. The operating system or boot kernel 201 begins the process. Initially there is some scheduling overhead 202. The single processor then executes processes 1, 2, 3 and 4 in sequence. Proceeding from one process to the next process adds some task-swap overhead 203, 204 and 205. There is synchronization overhead 206 and then the application is complete at 207.

6

On the multiple central processing unit system 300, the application begins with operating system or boot kernel 301. Initially there is some scheduling overhead 302. The single processor then executes processes 1, 2, 3 and 4 in parallel. There is synchronization overhead 303 and then the application is complete at 304.

Adding additional central processing units to execute parallel processes, however, does not come without risk. Parallelism is now found at the software process level, independent of the level at which the application programmer interacts. Writing an application for such a parallel system is much like writing an application for a single processor case. The application programmer is not concerned about when code will be scheduled to run. Only the operating system or boot-code scheduler takes that into account. This is a major benefit, since it is as easy to create applications on such a system as a single processor-based system and higher system performance is realizable. Only the operating system or boot-code programmer needs to be aware of the hardware. However this also presents unique challenges, since the application programmer is normally unaware of the hardware and the system must execute such code just as if it was made up of a single processor.

Any data shared between processes must be kept coherent. As a result, the software-processed based multiprocessor is less flexible in hardware than an application-division based multiprocessor. However, development on such a system is much easier and more scalable, allowing for greater degrees of parallelism and hence higher performance.

When hardware runs a software-process based multiprocessing solution it is required to keep shared data coherent. The application software written for a single-processor system must run correctly on a multiprocessor system. Through the use of symmetric multiprocessing (SMP), it is possible to satisfy both of these conditions. Conventional symmetric multiprocessing systems CSMP are commonly employed on desktop PCs (dual central processing units) and small-scale (4-way or 8-way) servers. Many of the same techniques can be used in an embedded application, but can be enhanced further by the tighter integration of an embedded central processing unit. Additional hardware can be employed to allow for better real-time determinism such as interrupts and peripheral management.

Symmetric multiprocessing derives its name from the premise that each central processing unit in the system behaves exactly the same as any another. All central processing units run the same instruction set, at the same frequency and all have access to all system resources. This is needed, because applications are written as if they are to be run on a single central processing unit. As a result, all central processing units that can run a process need to appear identical.

One of the greatest challenges to an SMP system is in keeping data coherent. Since the operating system or boot-code scheduler will launch different processes on each processor, any data that is used by more than one process must be kept current. A central processing unit that changes a shared variable must have that change reflected in the other central processing units. This may be done by having a large shared memory. By definition such a large shared memory does not allow for any local data memory. For performance and data coherence reasons, a data cache must also be employed when the base central processing unit instruction set does not support multiprocessing.

The embedded symmetric multiprocessing ESMP architecture of this invention will not have any serial message

passing. All data is kept in the shared memory and information is passed between processes in the form of shared variables. This is just the same as in the single-processor case, where the central processing unit will access the same memory locations for shared data between processes. However, in a multiprocessor model, shared-resources can be a bottleneck since only one central processing unit can have access to the data at a given time.

The greatest challenge from a data integrity viewpoint is making sure central processing unit registers are updated with any changes to the shared variables that may be stored. This is most conveniently done using good documentation and disciplined programming habits, declaring any variable or pointer that can be changed as a volatile type. This will force the central processing unit to load from main memory a new value into the register file any time the variable is used. However, since this is not a requirement in the single-processor case, it will cause a slight burden to the end application programmer in directly porting existing code. Changes to the compiler can also guarantee register coherence, since the compiler can generate code that will always reload data from main memory.

FIG. 4 illustrates a first generation 2-way embedded symmetric multi-processor ESMP architecture. A single flash memory 400 stores a single program stream. Both central processing units 401 and 403 receive their instructions from flash memory 400 via instruction bus 407 and program access and arbitration logic block 402. When an instruction cache miss occurs, arbitration logic 402 determines which processor has priority access to the flash memory 400. Both central processing units 401 and 403 receive their data likewise from the same source, internal shared data memory 404. All system resources are shared and visible to central processing units 401 and 403. Both central processing units 401 and 403 run the same instruction set and have identical organizations. Similarly, system peripherals and arbitration logic 406 is shared by both central processing units 401 and 403. Central processing unit 401 interacts with internal shared data memory 404 and systems peripheral arbitration logic block 406 via 32-bit data access bus 408. Central processing unit 402 interacts with internal shared data memory 404 and systems peripheral arbitration logic block 406 via 32-bit data access bus 409.

As illustrated in FIG. 4, program instructions are kept in a single external flash memory device 400. Alternately the instructions may be stored in an internal ROM, not shown. This is the same as the single-processor model. Since there is only one path to the instruction memory and each central processing unit 401 or 403 needs to access the program memory on nearly every cycle, the processors require an instruction cache for cycle performance reasons. This differs somewhat than a single-processor case, where the instruction cache is used due to variations in memory speed. Even if all of the program is kept internal to the device, an instruction cache near each central processing unit is needed. This prevents a performance bottleneck from occurring during program memory access arbitration. When both central processing units 401 and 403 suffer a cache miss, program access arbitration logic 402 will stall central processing unit based on hardware priority while the other central processing unit refills its cache line.

In ESMP there is symmetry between all the processors in the system. The processors are of the same type and, in normal operation, run at the same frequency. With this large amount of complex hardware operating at very high clock frequency, however, there is significant concern about power dissipation. Common techniques used to reduce power in

VLSI devices: (1) disable clocking on portions of the device that lay idle; and (2) reduce clock frequency in portions of the circuit that are in active operation but are doing minimal work.

FIG. 5 illustrates a prior art ESMP system. In FIG. 5 all devices are driven by the same clock to aid in coherence and software synchronization. System clock 510 is distributed from the clock controller 500 to two central processing units 503 and 504 and to memory and peripherals 506. Great pains are taken to reduce unwanted clock skew between blocks to a minimum. The system clock 510 is buffered with great care as it drives central processing units 503 and 504, and memory and peripherals 506.

In order to reduce power dissipation, it is desirable to have the option of allowing selected system central processing units run at lower frequencies if they are not being fully loaded. There are two main issues that must be resolved to support this mixed frequency operation. First, the performance penalty resulting from lost cycles associated with clock synchronization bears considerable investigation. Of even greater importance, however, is that any software techniques used by the operating system or boot-kernel to maintain process synchronization particularly at program barriers might need to be re-worked to accommodate the operation of selected central processing units are running at different speeds.

FIG. 6 illustrates the variable ESMP clocking system of this invention in high-level block diagram form. CPU-1 clock 601 and CPU-2 clock 602 are independently controlled by clock controller 610. Clock controller 610 receives clock source 600 and system select signals 609. Clock synchronizer/wait state generator block 605 selects either the system clock 603, CPU-1 clock 601 or CPU-2 clock 602 to become memory and peripheral clock 604 supplied to shared memory and peripherals 606. In FIG. 6, system clock 603 is the standard high frequency clock. CPU-1 clock 601 and CPU-2 clock 602 are possibly lower frequency, phase locked derivatives of system clock 603. Clock synchronizer/wait state generator block 605 controls the wait states of central processing unit 607 and 608 via wait state line 612.

FIG. 7 illustrates clock controller block 610 and system synchronization block 611 in more detail. Clock controller block 610 includes clock buffer 701, CPU clock select inputs block 702, phase-locked loop block 705 and frequency dividers 706. System synchronization block 611 includes clock control register 711, clock synchronizer/wait signal generator 710 and output multiplexer 720 for the memory-peripheral clock 604. System select signals 609 supplies configuration input bits. Clock source 600 supplies a stable clock source.

Assume first that normal clocking is used. In that case all clocks are identical. CPU-1 clock 601, CPU-2 clock 602 and memory-peripherals clock 604 are all the same frequency, but not necessarily the same as the source frequency clock 600. Data bit 703 is asserted as logical 1. Clock register 711 captures on the next clock the condition that all clocks are identical. This is applied via signal 712 to both clock synchronizer 710 and output multiplexer 720. Inside clock controller 610, a single phase-locked loop circuit 705 is used to generate CPU-1 clock 601, CPU-2 clock 602. At the output of system synchronization block 611 CPU-1 clock 601, CPU-2 clock 602 and memory-peripherals clock 604 would all have the same divide ratio with respect to the source clock. This would be a power of 2. Clock synchronization/wait state generator block 710 would use a bypass

mode. All clocks are the same, no synchronization takes place and no wait-state is needed.

If variable clocking is used, one of the CPU clocks, for example CPU-1 clock **601**, is divided down to a lower frequency. The following occurs. Memory-peripherals clock **604** is set to the slowest CPU-1 clock frequency. Software does this by the same system select input signal **609** that sets the CPU-1 clock to a lower frequency. The slowest CPU-1 clock must be within a power of 2 of the fastest CPU-2 clock. This assures that the system synchronization logic will work correctly.

A memory or peripheral access request by central processing unit **607**, employing the slowest clock in this example, proceeds as follows. Central processing unit **607** makes a memory/peripheral access request. Clock synchronizer/wait signal generator block **710** immediately creates a wait-state signal **612** back to central processing unit **607** delaying the request for one wait state. This one wait state assures that regardless of which central processing unit made the system request, the central processing unit will remain the same while the slower clock for the memory-peripherals first drives the logic for either a system read or write. On a write, the data bus will hold the data during this system wait cycle. On a read, data is placed on the data read bus, and when the wait state is de-asserted, the data will be latched into the central processing unit logic on the following central processing unit clock cycle. This approach assures data integrity. This one-wait state is effectively $\frac{1}{2}$ the period of the slower system clock. As a result, system synchronization block **611** either passes the system clock through when the central processing unit clocks are the same, or divides it by 2 and generates a wait state when operating in variable clock mode.

System clock **718** from clock synchronizer/wait state generator block **710** is a separate clock at the frequency of clock source **600**. This mechanism works properly because software sets memory-peripherals clock **604** at the slowest central processing unit clock. Control registers are not used inside system synchronization block **611** to control clocking during the memory/peripheral access because errors would result. These controls are kept in hardware and the only programmable feature inside system synchronization block **611** is selection of the bypass mode or the variable clocking mode.

In summary, if the system scheduler determines in software that one or more central processing unit is not fully loaded, then that central processing unit becomes a candidate for clock frequency reduction. Before any memory or peripheral accesses occur, the system clock must be synchronized with the clock of the specific central processing unit making the system access. Synchronization removes the possibility of malfunctions in the time interval during which

switching of clocks from a higher to a lower frequency and conversely from a lower to a higher frequency takes place. This kind of synchronization between phase locked signals of different frequency is relatively easy to achieve. The variable clocking system provides significant power reduction in the frequently occurring scenario where processors are not driven to their limits by temporary system requirements.

What is claimed is:

1. A data processing system comprising:
 - an input reference clock signal;
 - a plurality of central processing units;
 - at least one peripheral device connected to each of said plurality of central processing units;
 - a clock distribution system generating a plurality of clock frequency signals, supplying an independently selected one of said plurality of clock frequency signals to each central processing unit and supplying a lowest clock frequency signal supplied to any central processing unit to said at least one peripheral device.
2. The data processing system of claim 1, wherein:
 - each of said plurality of clock frequency signals are derived by dividing said input reference clock signal by an integer ratio.
3. The data processing system of claim 2, wherein:
 - the ratio between any two clock frequency signals is an integral power of two.
4. The data processing system of claim 1, wherein:
 - said clock distribution system independently assigns a wait state to each central processing unit to allow for synchronization with said at least one peripheral device whenever said clock frequency signal supplied to said at least one peripheral is less than a highest frequency clock used.
5. The data processing system of claim 1, further comprising:
 - a task scheduler for distributing tasks among said plurality of central processing units, said task scheduler causing said clock distribution system to independently select for a central processing unit one of said plurality of clock frequency signals less than a frequency of a highest frequency clock of said clock distribution system when said selected central processing unit is less than maximally loaded.
6. The data processing system of claim 1, wherein:
 - each of said plurality of central processing units includes a corresponding instruction cache; and
 - said clock distribution system supplies each said instruction cache with a clock frequency signal supplied to said corresponding central processing unit.

* * * * *